CoCoMANS: Computational Co-Design for Multi-scale Applications in the Natural Sciences

Principle Investigator:
Dana Knoll, T-3;
nol@lanl.gov

Co-PI:
Allen L. McPherson
CCS-7
mcpherson@lanl.gov

Co-PI:
Patrick S. McCormick
CCS-7;
pat@lanl.gov

Abstract

The advent of exascale computing resources presents both opportunities and challenges, and there is significant discussion about how to best utilize exascale computing resources. In the CoCoMANS project, we are addressing these challenges by forging a qualitatively new predictive-science capability that exploits evolving high-performance computer architectures for multiple application areas–including materials, plasmas, and climate–by simultaneously evolving the four corners of science, methods, software, and hardware in an integrated computational co-design process. We are developing new "applications-based," self-consistent, two-way, scale-bridging methods that have broad applicability to the targeted science, will map well to emerging heterogeneous computing models (while concurrently guiding hardware and software to maturity), and provide the algorithmic acceleration necessary to probe new scientific challenges at unprecedented scales.

Project Goals

The primary goal of this project is to contribute to the foundation for computational co-design at LANL and beyond. The CoCoMANS project will demonstrate a paradigm shift in at least two application areas on emerging architectures. By considering multiple application areas we not only capture a broader perspective of the co-design space but also promote the cross-fertilization of ideas. We are not proposing to solve a larger version of the same physics problem, utilizing the same algorithms, on a new computer. We are stepping back and asking: *What new physics problems, or physics problem formulations, can we solve at exascale, and how can we co-design numerical algorithms and their supporting implementations to achieve this innovation?* Additionally, CoCoMANS is joining forces with AMD, HP, Intel and NVIDIA to test and characterize current and emerging processors and node-level architectures, and employ hardware simulators and performance modeling to explore potential system designs with an eye towards the exascale era.

Co-Design Process

Our proposed co-design process is one of empirical test- and benchmark-driven experimentation. We will develop and implement an experimental framework, along with supporting documentation and tracking capabilities, to evaluate and test multiple numerical methods, along with evolving implementations of four types of codes: micro-benchmarks, kernels, mini-apps, and compact-apps. Our goal is to develop codes that approach an optimal implementation for some available candidate hardware and to identify bottlenecks and architectural shortcomings that can be addressed in future hardware designs. Where needed, this approach is strengthened by modeling and analysis. We will employ an established performance modeling methodology to characterize components of an implementation and analyze its performance on current or future architectures.

All of our implementations are built on the base of the numerical methods, solvers, and algorithms necessary to construct the required multi-scale applications. We will use existing prototyping software and well-understood development tools to experiment with numerical algorithms. These implementations are for analysis purposes only, and are not designed to run efficiently on emerging architectures—they act as abstract input to the various high-performance implementations.

Project personnel are working in four main functional areas (Fig. 1), with some important overlap. In domain science the tasks are to build applications that generate "real" scientific results. Solvers and numerical algorithms will be devised to build domain science applications. The computer science is concerned with implementing the applications, solvers, and numerical algorithms in the most efficient manner on a variety of hardware platforms. Hardware represents performance modeling, the interactions with external hardware partners, and physical equipment. In practice these domains are a continuum, and the expertise of project personnel span the nominal boundaries.

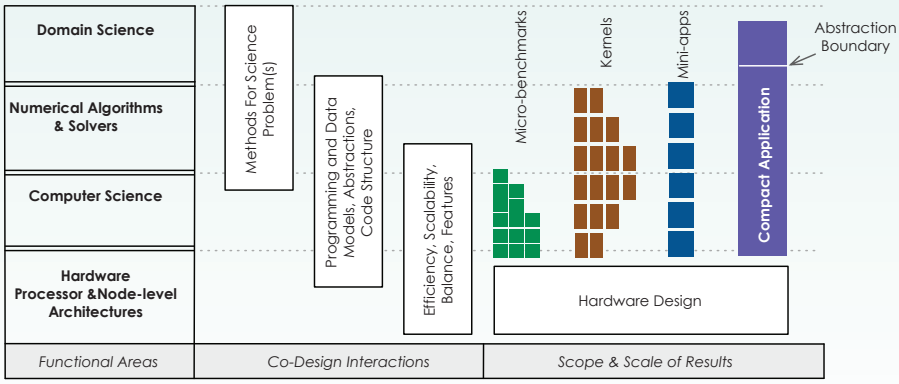

Fig. 1

We are developing a pragmatic process of progressive refinement based on co-design

interactions between these functional areas. Domain scientists will interact with numerical methods experts to develop optimal multi-scale physics algorithms for each domain. Computer scientists will implement these methods on various hardware configurations, run tests and experiments, providing results and feedback both back up the chain and to the hardware partners. While co-design interactions naturally focus on "nearest neighbor" relationships, the team is small enough that the "big picture" can be maintained.

Scale-bridging Algorithms

We build on foundations of algorithmic scale-bridging ideas that reside in statistical mechanics, kinetic theory, and transport theory. These ideas have matured in steady-state neutron transport for nuclear reactors. There, using this algorithm, the neutron transport equation would be referred to as the high order (HO) problem and is six-dimensional (6D). The low order (LO) problem is constructed from combining the 0^{th} and 1^{st} phase-space moments of the transport equation. This LO equation is three-dimensional (3D), and thus is a reduced-space representation of the HO problem. This LO (moment-based) problem is used to accelerate the convergence of a brute force transport source iteration via bridging the transport and diffusion scales. This algorithm has characteristics of a multigrid method. The LO problem lives in a reduced dimensional space (3D). It is an efficient way to relax the system-scale "diffusion" features that naturally reside in the HO problem. The more expensive HO problem (6D) is only required to relax the fine-scale transport features. The HO and LO problems are connected with restriction and prolongation processes related to the moments. A convergence plot comparing the brute force solution to the HO kinetic problem with the HO/LO scale-bridging algorithm on a typical model problem is shown in Fig. 2. This impressive convergence acceleration forms the basis of an exceptionally powerful scale-bridging algorithm, which has a very broad application potential and a number of characteristics that will be beneficial in moving to exascale computing. This general moment-based HO/LO algorithm has natural heterogeneity and concurrency. When statistical methods are used for the HO solver, the algorithm will also have some natural fault tolerance.

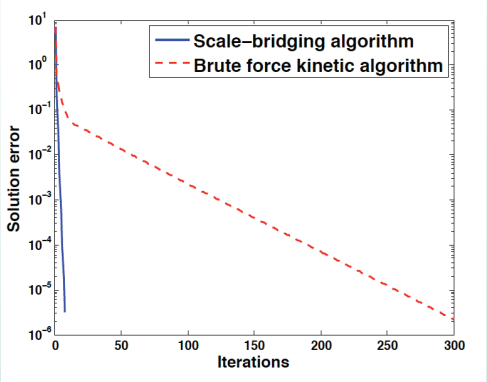

Fig. 2

Applications

Here we isolate the application area of kinetic plasma simulation as one specific example of a possible demonstration of a paradigm shift. We consider a historical magnetic reconnection problem: island coalescence. The physics of this problem is an intimate interaction between the system scale, more refined scales, and non-Maxwellian electron effects. Fig. 3 depicts this concept on the well-studied problem of magnetic island coalescence. Traditional particle-in-cell (PIC) algorithms must respect all electron scales over the entire geometry. Such algorithms have very restrictive constraints on grid spacing and time-step size. A recent, very impressive, PIC simulation of this problem used 20,480 cores on the Pleaides cluster (NASA), and required a spatial grid of 17920 x 8690 cells and 600,000 time steps (Fig. 3). Additionally, this simulation was forced to use an artificial ion/electron mass ratio of 25 (true value is 1836 for hydrogen plasma). Our new heterogeneous scale-bridging algorithm will enable the use of the true mass ratio (more faithful physics), less mesh by a factor of 4 in each direction, and time steps that are 100 times larger, and thus a factor of 100 improvement in simulation effort. We expect this algorithm to map very effectively to advanced architectures. Thus, the total computational co-design impact on this particular simulation could range from a factor of 200- to 1000-fold improvement. Given appropriate algorithm maturity and mini-app support, we will use this problem to demonstrate our paradigm shift in kinetic plasma simulation on emerging architectures. With this paradigm shift demonstrated, we will have defined a path to a 3D system-scale kinetic plasma simulation that is not currently possible. Similar paradigm-shifting studies with compact-apps in material science and climate are planned, but we will not elaborate here due to space limitations.

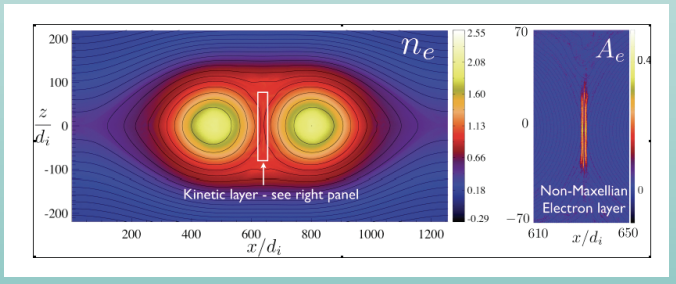

Fig. 3

This project is supported by LANL LDRD